

---

# NASA's Supercomputing Experience

---

F. Ron Bailey

---

(NASA-TM-102890) NASA'S SUPERCOMPUTING  
EXPERIENCE (NASA) 30 p CSCL 12A

N91-14727

Unclas  
G3/59 0321369

December 1990



National Aeronautics and  
Space Administration



---

# NASA's Supercomputing Experience

---

F. Ron Bailey, Ames Research Center, Moffett Field, California

December 1990



National Aeronautics and  
Space Administration

**Ames Research Center**  
Moffett Field, California 94035-1000



# **NASA's SUPERCOMPUTING EXPERIENCE**

**F. Ron Bailey**  
**NASA Ames Research Center**  
**Moffett Field, CA 94035, USA**

## **ABSTRACT**

A brief overview of NASA's recent experience in supercomputing is presented from two perspectives: early systems development and advanced supercomputing applications. NASA's role in supercomputing systems development is illustrated by discussion of activities carried out by the Numerical Aerodynamic Simulation Program. Current capabilities in advanced technology applications are illustrated with examples in turbulence physics, aerodynamics, aerothermodynamics, chemistry, and structural mechanics. Capabilities in science applications are illustrated by examples in astrophysics and atmospheric modeling. The paper concludes with a brief comment on the future directions and NASA's new High Performance Computing Program.

## **1. Introduction**

The National Aeronautics and Space Administration (NASA) has for many years been one of the pioneers in the development of supercomputing systems and in the application of supercomputers to solve challenging problems in science and engineering. Today supercomputing is an essential tool within the Agency. There are supercomputer installations at every NASA research and development center and their use has enabled NASA to develop new aerospace technologies and to make new scientific discoveries that could not have been accomplished without them.

NASA's experience with advanced supercomputers began in the 1970s with the installation of the Illiac-IV parallel computer at Ames Research Center and the CDC STAR-100 vector computer at Langley Research Center. These two systems were the first of a new generation of high-performance computers designed to achieve dramatic performance increases through the architectural innovations of parallel and vector processing. Although these systems were marginally successful, they provided a valuable testbed environment in which to learn to effectively use parallel and vector computing techniques. Architectural investigations continued into the 1980s with the development of the Massively Parallel Processor, a massively parallel SIMD processor system, at Goddard Space Flight Center, and the Cosmic Cube/Hypercube, parallel MIMD systems, at the California Institute of Technology and NASA's Jet Propulsion Laboratory. In the 1980s NASA also established the Numerical Aerodynamic Simulation (NAS) Program to pioneer supercomputing technologies and make them available to the aerospace community. The NAS Program, for example, acquired the first full-scale Cray-2 and first Cray YMP computer systems and developed the first supercomputer UNIX environment.

NASA's involvement in supercomputing applications also began in the 1970s. Acquisition of the Illiac-IV, STAR-100, and Cray-1 computers spurred the development of innovative parallel computing algorithms. New computational disciplines, such as computational fluid dynamics and computational chemistry soon emerged and it became clear that supercomputers were a truly enabling technology. Today, supercomputers are being applied routinely to advance technologies in aeronautics, transatmospherics, space transportation, and space exploration, and to study the Earth, other planetary systems, and objects in outer space.

The purpose of this paper is to provide a brief overview of selected NASA activities in supercomputing from two perspectives: supercomputing system development and supercomputing applications. The NAS Program is used as an example of NASA's effort to remain at the leading edge of supercomputing system technology and bring this technology to the aerospace research and development community. NASA's use of supercomputers in advanced technology applications is illustrated through examples in turbulence physics, aerodynamics, aerothermodynamics, chemistry, and structural mechanics. Examples in atmospheric science and astrophysics are presented to illustrate supercomputing applications in science. Finally, the paper concludes with a look into future directions in high-performance computing.

## **2. Numerical Aerodynamic Simulation Program**

The NAS Program<sup>1</sup> has two major goals. The first goal is to provide a national computational capability that is available to NASA, Department of Defense (DOD), other government agencies, industry, and universities, as a necessary element in ensuring continued leadership in computational aerodynamics and related disciplines. The second goal is to act as a pathfinder in advanced large-scale computing capability through systematic incorporation of state-of-the-art improvements in computer hardware and software technologies. The NAS Program began full operation in 1986 and today provides service to over 1700 users at 133 locations in the United States.

The NAS complex of computers, called the NAS Processing System Network (NPSN), is shown in Fig. 1. It currently includes Cray-2 and Cray YMP/8128 supercomputers, an Amdahl 5880 mainframe computer and associated mass storage system with 2.3 terabytes of on-line capacity, 5 VAX minicomputers, 80 Silicon Graphics workstations, a 32,000-processor Connection Machine parallel computer, and a 128-processor Intel Touchstone Gamma prototype parallel computer. All computer systems are linked through an extensive local area network combining Ethernet, HYPERchannel, and Ultra-net technologies to provide hardware data rates ranging from 10 to 800 mbits/sec. All computers operate under the UNIX operating system with DOD internet (TCP/IP) network communications provided via the well-known Berkeley UNIX "r" commands. Thus, a user can access any computer, run jobs, and transfer data among computers using

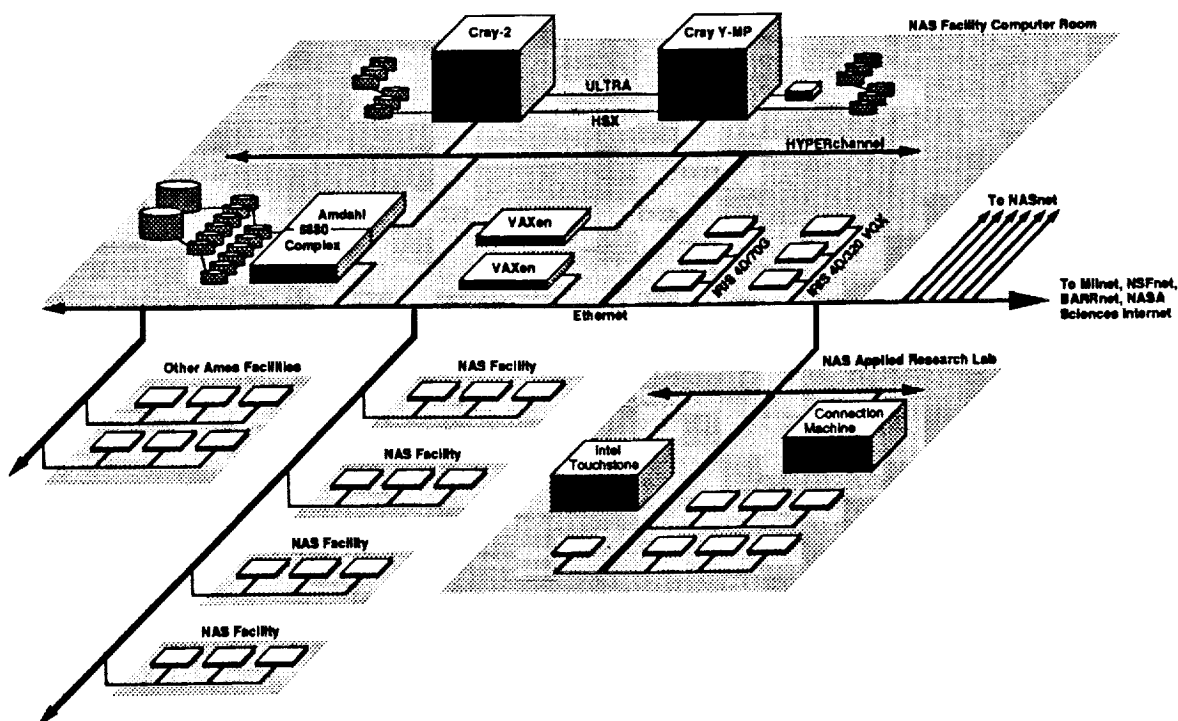


Fig. 1 NAS Processing System Network (NPSN).

a single set of commands on all computers. A more complete description of the NPSN is given in references 2 and 3.

Remote users access the NPSN through a wide choice of existing national communication networks including the DOD-sponsored MILnet, the National Science Foundation-sponsored NSFnet, and NASA's Science Internet. NAS is also a member of the Bay Area Regional Research Network that provides 1544 kbits/sec links between Ames Research Center, Bay Area universities, Department of Energy laboratories, and industry. MILnet is the main communication path to DOD laboratories, while NSFnet, and BARRnet serve the university community. For NASA and aerospace industry users, NAS has developed NASnet. NASnet uses Ethernet bridging technology to provide fast response and high throughput to remote graphics workstation users. NASnet currently serves 31 remote sites at bandwidths ranging from 56 to 1544 kbits/sec.

The NAS Program's experiences in advanced supercomputer technology began in 1984 when it decided to acquire Cray Research, Inc.'s Cray-2 computer, which was in design at the time. NAS, in collaboration with the Department of Energy, began prototype testing in 1985. In late 1985, NAS took delivery of the first full-scale Cray-2 configuration. This system had a memory capacity of 2.048 gigabytes (over 268 million 64-bit words) which was 34 times the central memory capacity of the Cray XMP/48. The Cray-2's very large memory, coupled with its ability to sustain a computing rate of 250 million

floating-point operations per second, allowed computational fluid dynamics researchers to increase their application program size by more than an order of a magnitude and increase their complexity without significant code modification. The very large memory virtually eliminated memory-disk transfer overhead during execution of large programs. The net effect was decreased preparation time for researchers tackling increasingly complex problems, plus improved job turnaround time. More importantly, researchers were able to perform fluid dynamics simulations about realistic three-dimensional (3-D) aerospace vehicle configurations that were not possible just a few years before.

Providing a uniform user interface and operating environment across subsystems was a major NPSN design goal. Furthermore, the environment had to remain common across subsystems even as new supercomputer generations and other new subsystems were added. Only an open system architecture (implementation based on openly available definitions for hardware and software interfaces) that was free of vendor proprietary restrictions could meet this goal. UNIX operation system, DOD internet protocols, and the Berkeley network interface were chosen and implemented on the Cray-2. Prior to this, UNIX had only been implemented on minicomputers and workstations, and many felt that UNIX was not suitable for supercomputers. Today, UNIX is the standard for supercomputer operating systems. The key to the successful implementation on the Cray-2 and later on other supercomputers was to treat UNIX as a definition that could be implemented and, if necessary, enhanced to meet supercomputing needs. For example, a batch resource management system, NQS, was added to provide users the flexibility of both interactive and batch supercomputing. Moreover, UNIX, when coupled with common networking commands, creates a so-called seamless network environment that enables users to move easily between systems, allows for easy file access and command initiation across machine boundaries, and enhances user-code portability. Finally, the open architecture concept allows for modularity essential for ready implementation of new capabilities and has proved beneficial in the installation of the Cray YMP and other new systems.

Enhanced user productivity was another major NPSN design goal. To meet this goal, the NPSN was designed to include scientific workstations as the researcher's interface to supercomputers. The design was implemented by installing NAS Silicon Graphics IRIS workstations and connecting them directly to the Cray-2 through a local area network. This was a unique implementation at the time and has since been replicated at many supercomputer installations. Combined with UNIX and network software, the supercomputer becomes literally an extension of the workstation. The user resides at his workstation where text and small data files are created and modified, where applications are submitted and interacted with, and where complex input data and results are displayed and analyzed. Because numerical simulations produce an enormous amount of data, the workstation's graphics display capability is especially beneficial. By using graphics workstations, researchers can digest results in a reasonable amount of time and gain a qualitative picture of the physical phenomenon simulated. Numerous graphics packages



have been developed to operate across workstations and Cray supercomputers in a distributed manner and to allow researchers to display any facet of their data for easy interpretation and evaluation<sup>4-7</sup>. Today, interactive graphical post-data processing is an essential aspect of today's sophisticated numerical simulation process.

Providing ready access by the national aerospace research community from their home locations was another NPSN design goal. This led to the development of NASnet, a unique high-performance nationwide communication network that links remote user locations to the NAS facility via land-links, with communication speeds up to 1544 kbits/sec. NASnet connects the communication lines from the NAS facility to local area networks at each remote site through Ethernet bridges and thereby creates an extension of the NPSN local network. The implementation of UNIX and network software on remote workstations provides an environment in which remote researchers have virtually the same interactive capability as users at the NAS facility itself. Presently, NASnet is being redesigned to utilize new internet gateway technology and thereby increase its performance to approximately 6200 Kbits/sec.

### **3. Computational Fluid Dynamics**

Computational fluid dynamics has been one of the earliest drivers in advancing supercomputing within NASA. The introduction of numerical techniques in the late 1960s to treat highly nonlinear flows, such as transonic flow over lifting airfoils with embedded shock waves<sup>8</sup>, rapidly led to greatly increased interest in computational solutions to fluid flow problems. The discipline grew rapidly with improvements in both computer systems and algorithms. Fig. 2 shows that improvements in computer performance have been closely paralleled by improvements in numerical methods<sup>9</sup>. The data shown here illustrates that the cost of performing a given calculation has decreased three orders of magnitude due to advances in computers and also three orders of magnitude due to improvements in numerical methods over a 15-year period. While the example shown is for methods used to solve two approximating forms of the Navier-Stokes equations that govern the fluid dynamics of a perfect gas, it is indicative of the trend in other disciplines as well.

Computational fluid dynamics is now a well-established discipline within NASA and the aerospace industry<sup>10,11</sup>. While industry's main focus is on design applications, NASA is more involved in research. The objectives of NASA's efforts, which are discussed below, are to (a) advance fundamental understanding of basic fluid dynamic processes and (b) develop an advanced base of new aerodynamics technology for application to future generations of aerospace vehicles and propulsion systems.




Fig. 2 Comparison of numerical simulation cost trend resulting from improvements in fluid dynamics algorithms and in computers.

### 3.1 Turbulence Physics

Turbulence is the most common fluid motion encountered in nature. It greatly affects the resistance of bodies moving through fluids, fluid mixing, and heat transfer. Because of its importance, turbulence study has attracted some of the best scientific minds for over a century. Despite this attention, turbulence is still not well enough understood to be effectively controlled in engineering applications. If it could be controlled, the benefits would be significant. For example, in the case of a fleet of transport aircraft a 20 percent reduction in fuselage skin friction alone would save approximately \$1 billion a year in fuel consumption<sup>12</sup>.

With the development of supercomputers and highly accurate numerical algorithms, it is possible to obtain numerical solutions to the Navier-Stokes equations for simplified

geometry and at low Reynolds numbers. These numerical simulations are based on first principles and can be treated as though they were observed data from a physical experiment. Used in this way, numerical simulations are providing new insights into interrelationships between flow variables, such as the relationship between pressure and strain for which no measurement has yet been possible.

In addition to quantitative information of unprecedented detail, numerical solution offers, for the first time, the opportunity to visualize turbulent flows as they evolve in both space and time without the averaging and filtering inherent in results obtained from laboratory experiments. Each structure in a turbulent flow can be inspected prior to calculating the statistical characteristics of significant flow features. This visualization capability is valuable both for identifying new relationships among the various elements of the dominant flow structures and for suggesting new types of physical experiments. Fig. 3 shows the calculated vortical structures near a flat plate<sup>13</sup>. These results were obtained




Fig. 3 Calculated vortical flow structures in a turbulent boundary layer.

from the database provided by direct (first principles) numerical simulation of a flat-plate turbulent boundary layer<sup>14</sup>. The direct simulation provided pressure and all three components of the flow velocity for each time step and for each of 9.4 million grid points in the computational domain. A problem of this magnitude requires about 200 hours of CPU time on a Cray-2 computer. Turbulence physicists have analyzed the results shown in Fig. 3 and have concluded that locations of significant contributions to Reynolds stresses occur adjacent to large hook-shaped vortical structures that are common in turbulent flows<sup>13</sup>. Information of this sort aids in the development of statistical models that accurately reflect the underlying physical behavior of turbulence.

### 3.2 Aerodynamics

Computational aerodynamics is concerned with numerical simulations of aircraft flying at normal altitudes and at Mach numbers below five where air behaves as a perfect gas. The computational models employed differ in levels of approximation to the governing Navier-Stokes equations and in geometric complexity. Due to computer limitations, one generally is forced to trade geometric simplicity for less approximate solutions or more approximate solutions for geometric complexity<sup>9</sup>. As the approximations become less severe and the solutions become more complex, the computer resources needed for a solution grow rapidly<sup>15</sup>. Generally, approximations that neglect viscosity (inviscid) in the flow field and treat viscosity only in the boundary layer are in widespread use for aerodynamically clean configurations at or near cruise conditions. Formulations that include viscous terms become important when treating flows where the boundary layer is separated from the surface. These occur, for example, near boundaries of the flight envelope, in engine inlets and with those flows associated with vortex-enhanced lift.

The level of modeling receiving the most attention within the NASA research environment is Reynolds-Averaged Navier-Stokes (RANS). RANS models include all terms of the Navier-Stokes equations averaged over a scale that is large with respect to the turbulent eddy fluctuations, yet small relative to the macroscopic flow changes. This formulation introduces terms representing the transport of energy and momentum between the mean and fluctuating component. These cannot be determined analytically and must be phenomenologically modeled. Turbulence modeling is the major item pacing the full exploitation of the RANS approach and continues to be a focus of experimental and computational research.

Increases in computer performance have dramatically affected the utility of the RANS formulation<sup>16</sup>. Fig. 4 illustrates the effect of increasing supercomputer performance on advancing RANS simulation capability during the last 15 years. In 1975, a two-dimensional (2-D) calculation of transonic flow over a simple airfoil strained the capability of the CDC 7600 systems available at that time. The increased performance of the Cray-1 made it possible to compute simple 3-D wing flows by 1983. By 1986 it was

## INCREASES IN SUPERCOMPUTER POWER ENABLE ADVANCES IN COMPUTATIONAL FLUID DYNAMICS

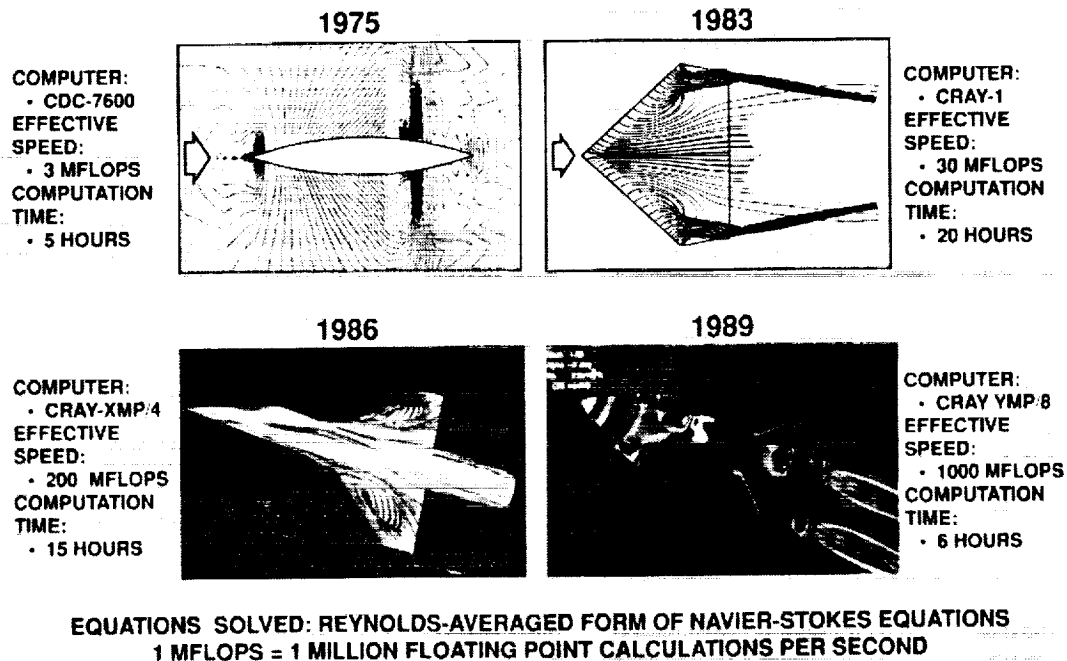


Fig. 4 Advances in Reynolds-Averaged Navier-Stokes simulation capability with increases in supercomputer performance.

possible to compute flows over wing bodies on the Cray XMP/4. Today the Cray YMP/8 enables simulation about very complex geometries such as the space shuttle configuration shown at the bottom right-hand side.

Fig. 5 shows more detailed results from the use of the RANS approximation to compute the transonic flow about the integrated space shuttle in ascent mode<sup>17</sup>. Here, computed surface pressures are compared with those obtained from wind tunnel tests. The computed and measured results are in excellent agreement over large portions of this complex configuration. A limited amount of flight test pressure data is available and is compared with wind tunnel and computed pressure along the side of the fuselage in Fig. 6. At this particular Mach number of 1.05, the computation is in better agreement with the flight data than with the wind tunnel data, apparently due to wind tunnel wall interference. Although there is disagreement found in localized regions, the computational approach is considered adequate to predict the aerodynamic loads needed to refine the current flight envelope and to evaluate proposed emergency abort maneuvers.

ORIGINAL PAGE IS  
OF POOR QUALITY




Fig. 5 Comparison of pressure contours between computation (outermost) and wind tunnel (innermost),  $M_\infty = 1.05$ ,  $\alpha = -3^\circ$ , and  $Re = 4.0 \times 10^6/ft$  (3% model). (a) Side view, (b) top view.

#### 4. Aerothermodynamics

Aerothermodynamics combines the discipline of computational fluid dynamics with the modeling of real gas properties. Vehicles flying at Mach numbers greater than approximately five create shock waves strong enough to cause the chemistry of the air to change significantly. Therefore, to adequately simulate such flows, additional equations governing the behavior of the real gas chemistry must be added to the Navier-Stokes equations. Clearly, the level of computational complexity increases over that of computational fluid dynamics alone.




Fig. 6 Comparison of pressure coefficient ( $C_p$ ) from computation (—), wind tunnel (o), and flight test ( $\nabla$  right side,  $\Delta$  left side) along  $\phi = 70^\circ$  line side of the orbiter fuselage,  $M_\infty = 1.05$ ,  $\alpha = -3^\circ$ ,  $Re = 4.0 \times 10^6/\text{ft}$  and  $10^\circ/9^\circ$  elevon deflection.

The goal of computational aerothermodynamics is to predict aerothermal environments both in the case of external flow around vehicles and the internal flow within propulsion units. The physical phenomena which must be addressed include aerodynamic forces and moments, convective and radiative heating rates, gas and surface interactions, catalytic surfaces, interactions with active cooling of Thermal Protection System (TPS) materials, and plasma layers and their effect on electronic communications to and from the vehicle. Understanding these phenomena is critical to the design of the next generation aerospace transportation systems and to ensure that they are adequately protected from the very high thermal loads encountered in hypervelocity flight. At the current state-of-the-art, the RANS equations coupled with the real-gas effects form the basis for the solutions to these phenomena.

Real-gas effects include radiative thermochemical nonequilibrium where finite-rate processes for chemical- and energy-exchange phenomena occur. To account for reactions, conservation equations for each chemical species must be added to the RANS equation set. For a relatively simple case of dissociating and ionizing air, one must typically consider nine air species. Including the associated conservation equations for each species nearly triples the number of equations to be solved. In addition, to account for thermal

nonequilibrium there are additional energy equations to describe the energy exchange between the various energy modes (translational, rotational, vibrational, and electronic). Finally, when thermal radiation is important, the radiative flux must be included in the energy equation.

Supercomputers have been essential in the development of computational aerothermodynamics and have permitted the modeling of complex 3-D flow fields with finite-rate chemistry. As an example, Fig. 7 shows the thermochemical nonequilibrium flow field about the Aeroassist Flight Experiment (AFE) vehicle at one of its trajectory points<sup>18</sup>. The left side of the figure shows computed normalized density contours. Superimposed on these are circles representing the location of the bow shock as determined from experimental shadowgraphs obtained from ballistic range tests. The experimental and computational bow shock locations are nearly identical. Shown on the right side of the figure is a comparison of the vibrational and translational stagnation line temperatures. As illustrated, there is a significant difference in the two temperatures. This difference, depending on the entry conditions, can result in a significant increase in the total radiative heating to the vehicle<sup>19</sup>.




Fig. 7 Thermochemical nonequilibrium flow field about the aeroassist flight experiment (AFE) vehicle. (a) Computed density contours, computed and experimental bow shock locations; (b) vibrational and translational temperatures along the stagnation line.



## 5. Computational Chemistry

During the past decade, computational chemistry has become a dynamic and versatile tool to assist in understanding the fundamental nature of a variety of substances, notably gases, materials, and their interactions. For example, computational chemistry methods are currently being used to accurately predict radiative intensity factors, high-temperature transport properties, reaction rates, and rates of energy exchange during particle collisions. These data are required by computational aerothermodynamics to simulate both external and internal flow fields of aerospace vehicles. In addition, computational chemistry methods are being used to study and design space-based materials which are more resistant to the harsh environment of space and, in particular, the impact of energetic oxygen atoms.

Computational chemistry involves solving the Schrodinger equation to determine the properties of atoms, molecules, and matter. The techniques used to solve this equation are well-documented and will not be described here. It should be noted, however, that supercomputers have had a profound influence on computational chemistry. While the performance gains of an order of magnitude have naturally allowed much larger problems to be considered, supercomputers have had a far greater impact on computational techniques. For example, the ability to perform Full Configuration Interaction (FCI) calculations in realistic basis sets has provided researchers with detailed benchmarks for evaluation of more approximate methods. These FCI calculations<sup>20</sup>, using the Cray-2 and Cray YMP, have produced results that are superior in accuracy to those determined from experimental techniques. Furthermore, these results have demonstrated that multireference CI wave functions reproduce the FCI results to very high accuracy and use approximately one-fourth of the computer time. Finally, the speed of supercomputers has tremendously reduced the time required to obtain results. Previously, even if a desired accuracy could be obtained, the real time required to obtain results was frequently too long for the results to be useful or to impact a specific program.

High-temperature reaction rate constants for air and air plus hydrogen species are crucial to the accurate prediction of the external flow around hypersonic vehicles and modeling of the engine performance. Experimental rate constants are seldom available and almost impossible to obtain at high temperature. When data are available, they are often uncertain by several orders of magnitude. Fig. 8 is a comparison of the theoretical rate constant<sup>21</sup> for the three-body recombination of  $H + H + H_2$  with a variety of experimental data. As shown, at high temperatures the scatter in the experimental data is three orders of magnitude. Reaction rate constant calculations like these are providing data which do not exist elsewhere, as well as helping to clarify what is often a confusing maze of experimental data.




Fig. 8 Comparison of calculated and experimental rate constant for the three-body recombination of  $\text{H}+\text{H}+\text{H}^2$ .

## 6. Structural Mechanics

Computational structural mechanics is another major computational thrust within NASA and has been greatly influenced by supercomputers<sup>22</sup>. The finite element method, developed by the aerospace industry, is the predominant technique used for the analysis of solid and structural mechanics problems. In the early 1970s, NASA became heavily involved in finite element modeling through the development of the widely used NASTRAN finite element program. Today, finite element techniques employing supercomputers are widely used to solve linear and nonlinear problems for both static and dynamic analysis. Application areas include structural analysis of aerospace vehicles, propulsion systems, and space structures. The application of advanced composite materials to vehicle and propulsion system structures, in particular, has required the development of much more sophisticated numerical models that have in turn placed increased demands on computer capability.

An example of the application of supercomputers to a structural problem of critical importance is the space shuttle solid rocket booster (SRB) structural design assessment and redesign efforts following the accident that destroyed the Challenger. The accident

was believed to have been caused by the failure of a case joint in the right rocket motor<sup>23</sup>. An analysis of the joint was performed at three levels of detail<sup>24</sup> shown in Fig. 9. The first level consisted of a 2-D stress analysis for a shell model of the entire SRB that included an equivalent stiffness model of the joint. At the second level, the displacements calculated at level 1 were used as boundary conditions for a more refined 2-D shell model of an SRB segment including the joint. Finally, at the third level the displacement calculated at level 2 was used as a boundary condition for a detailed 3-D analysis of a 1-degree circumferential slice of the joint. Analysis was performed on both the original and redesigned joints. The results showed that the redesigned joint substantially reduced the gap motion at the O-ring surface and the sensitivity to O-ring performance.

The SRB analysis was performed using various NASA computer systems. The researchers were located at Langley Research Center and used a minicomputer for model preparation and verification and for post processing. The 2-D shell finite element structural analysis, using 54,870 equations, and the 3-D solid finite element analysis were performed on Cray supercomputers at Ames Research Center. This effort, performed in 1985 and 1986, was a graphic illustration of how effectively supercomputers can be used from far remote locations. The solution time for the full SRB model (level 1) exceeded 14 hours on a VAX 11/780 computer and was reduced to 14 minutes on a single Cray-2




Fig. 9 Three levels of modeling detail for solid rocket booster (SRB) case joint structural analysis.

processor when fully vectorized. More recent algorithm improvements have reduced the solution time to 13 seconds on the Cray-2 using 4 processors and microtasking and to 6 seconds on the Cray YMP using 8 processors and microtasking. In the latter case the 9.2 billion floating-point operations for matrix factorization took 6 seconds, a rate of 1.517 billion calculations per second.

## 7. Astrophysics

Supercomputers are also heavily used within NASA to perform astrophysical investigations. For example, the use of computational methods to investigate the formation and evolution of astronomical objects (stars, galaxies, radio sources, black holes, X-ray sources, interstellar clouds, etc.) has proven to be an essential part of the overall scientific approach to understanding the universe. Computation has affected nearly every area of astrophysics. It is a tool for interpreting observations that complements analytic theory and permits the study of complex phenomena that cannot be done in the laboratory.

Applications of supercomputers to astronomy is too broad a topic to cover here; thus, attention is focused on computer application of galactic formation and evolution<sup>25</sup>. Three features drive present research efforts and give rise to the need for the computational power of supercomputers. These are (a) the complex play of many physical processes requiring nonlinear interaction models, (b) greatly improved observational detail at all wavelengths requiring complex source modeling, and (c) reduction of raw observational data requiring high-speed reduction, enhancement, and analysis.

The study of the global dynamics of galaxies leads one to concentrate on the dynamics of the stellar population. The stars make up the bulk of the mass in most galaxies of interest. Stars move freely through these systems and they practically never collide. The collision mean free path is large compared to the system dimensions. The study of the stellar dynamics is then the study of a collisionless, gravitational n-body problem. The n-bodies are the stars, and they act like mass-points, or particles, that interact with the  $1/r^2$  forces of Newtonian gravitation. The models used for numerical experiments describe the time development of the self-gravitational particle system. Early experiments dealt with 2-D systems, but, with the advent of supercomputers in the late 1970s, fully 3-D experiments were made possible. The number of particles followed in an experiment is typically 100,000. The forces necessary to integrate the equations of motion of the particles are obtained from potentials that are calculated on a grid. These potentials are calculated from the tabulated star densities by solving the Poisson equation at each integration step and the equations of motion are integrated by the time-centered leapfrog method.

The approach described above has been used to study galaxy formulation and internal dynamics and the interaction of galaxies<sup>26</sup>. One example is the investigation of star formation stimulated by the collision of galaxies. Far infrared measurements from the

Infrared Astronomy Satellite have revealed a class of very luminous interacting galaxies which appear to be starburst galaxies. These results have stimulated research on the nature of the interactions and the ways in which the star-formation process can be stimulated through the interaction. An example of an optical photograph of an interacting pair of galaxies is shown in Fig. 10. The photograph was made in such a way to emphasize the regions of high star formation. The observations show vividly the violent disruption of the two galaxies and the regions of high current star formation, particularly at the centers of the two galaxies.

One of the computational approaches to understanding the starburst phenomenon lies in understanding the dynamical interaction of two spiral galaxies embedded in massive halos by numerical experiments. For example, one can follow the dynamical interaction of two disk galaxies to understand the nature of the contraction and subsequent expansion of the two galaxies. Fig. 11 shows four snapshots of a collision sequence in a numerical




Fig. 10 Interacting galaxies UGC 12914/12915. (Photograph courtesy of Dr. H. Bushouse, NASA Ames Research Center.)




Fig. 11 Time development of numerical experiment on interacting galaxies as viewed from above the orbital plane.

experiment<sup>26</sup>. The disks were oriented in the plane of the orbit and both the disk particles and the surrounding halo particles are shown. The particles shown are a sampling (2048) of the total number of particles in the calculation (100,000). The counter in the bottom right of the frames is the elapsed time in nondimensional units. The time sequence shown corresponds to roughly 150 million years. Understanding the nature of the interaction and the time variations of the densities and momenta in the two galaxies is the first step in a longer term project toward understanding the nature of the star-formation process in these systems.

## 8. Atmospheric Modeling

Atmospheric modeling on supercomputers has become a principal tool in predicting weather and studying climate change. The importance of this activity is clearly evident in light of the need to assess the influence of man's activity on the environment. NASA has

been a major contributor to advancing the applications of supercomputers to this field including: global modeling of the interactions between the oceans, atmosphere, land and ice; simulation of the dynamics and chemistry of the Earth's atmosphere including the response of the Earth's ozone layer to changing chemical atmospheric composition; and assimilation of data acquired from spaceborne sensing instruments. NASA is also engaged in studying the atmospheres of the planets. The study not only increases our understanding of the planets, but also lets us gain insight into the dynamics of the Earth's atmosphere.

The modeling of Mars' water cycle is one example of the use of supercomputers to study planetary atmospheres. Spacecraft missions to Mars during the 1960s and '70s have shown that the total amount of water vapor in its atmosphere fluctuates during the course of a year by an amount equivalent to about 1 km<sup>3</sup> of ice. This indicates that Mars has an active hydrological cycle in which water is exchanged between the atmosphere and various other reservoirs. The nature and distribution of these other reservoirs is not entirely clear, but it is known that one of them is the "residual" north polar cap—not the one that waxes and wanes with the seasons because that one is made of CO<sub>2</sub> frost—but the one that remains throughout the summer. In other words, when the seasonal CO<sub>2</sub> ice cap finally sublimates away in late spring, it leaves behind an underlying water ice cap that is less volatile and can survive the summer's heat. This remnant polar cap serves as a source for atmospheric water vapor. Surprisingly, however, the seasonal CO<sub>2</sub> cap at the south pole never disappears. For reasons not understood, carbon dioxide survives at the south pole all year long.

The significance of this asymmetry in polar cap behavior is that: if the north cap supplies water to the atmosphere each summer, and the south cap does not, then where does the water go? There have been many suggestions, but all of them are based on models that treat atmospheric transport so crudely that the results cannot be considered proof. It is this aspect of the problem, atmospheric transport, that requires supercomputers. To model atmospheric transport properly, one needs to solve the full 3-D equations of motion on a globe, and carry out the calculations for very long periods of time. The temporal aspect of the problem is important in order to know net annual gains or losses. Thus, multi-year simulations are required (one Mars year is about two Earth years). Even this, however, may not be enough to properly model the problem. The reason is that in addition to transport, there are important physical processes that occur, such as cloud formation, precipitation, boundary layer mixing, etc., that can affect the behavior of atmospheric water just as much as transport. Adding "physics" to the problem could easily increase the computational demand by an order of magnitude.

During the past several years, NASA researchers at Ames Research Center have been using the Cray supercomputers at NAS to attack these problems. They have constructed a transport model as well as various physical models. The speed and memory of these supercomputers have enabled the examination of polar mixing processes, cloud

formation, and boundary layer exchange at a level of detail not previously possible<sup>27</sup>. For example, the ability of its atmosphere to mix water vapor into the polar regions during winter is a key issue concerning the current Mars climate system. In Fig. 12 contours of a nearly conserved dynamical quantity (isentropic potential vorticity, IPV) are shown to illustrate this mixing process. The IPV field is taken from a simulation of the Martian winter circulation using a 3-D climate model being developed by researchers at Ames Research Center and based on the work in reference 27. The center of the figure is the north pole while the edges are at about 30 degrees north latitude. The large gradients in




Fig. 12 Simulated isentropic potential vorticity (IPV) field at the Martian north pole in winter.




the IPV field just poleward of the dark ribbon surrounding the pole represent a barrier that large-scale wave motions must break down and penetrate if mixing is to occur. At this particular time a wave number of three disturbance is evident, but its amplitude is small and the mixing it produces is weak. For this particular experiment, therefore, the Martian polar vortex is relatively stable and similar in this respect to the Earth's Antarctic stratosphere during winter. By running the transport code at high temporal and spatial resolution, for example, it was found that the wintertime circulation on Mars was sluggish in its ability to move water to the very high latitudes, which is contrary to earlier suggestions. Similarly, we have found that when clouds do form on Mars, they do not readily precipitate water to the ground; this also is contrary to earlier suggestions.

Recently supercomputers have been used to solve one of the most perplexing problems of atmosphere-surface interactions on Mars<sup>28</sup>. Namely, how does water move back into the surface? All previous studies indicate that once water is driven out of the soil, it does not readily return. By running a coupled boundary layer/soil model at very high vertical and temporal resolution, a physical process has been discovered that facilitates the return. This will significantly affect the field and open up new areas of research. Similar statements can be made about recent transport and cloud studies.

## **9. A Look to the Future**

To meet future challenges in computational science, NASA has initiated the High Performance Computer (HPC) Program. The goal of the HPC Program is to accelerate the development and application of high-performance computing technologies to meet science and engineering requirements for continued leadership in aerospace. Timely realization of the goals in many of NASA's science and engineering missions require significantly increased computing power. For example, in the field of computational fluid dynamics, the governing equations have been known for over a century, but, it has required supercomputer-level performance to develop tools for understanding the physics of aerodynamic flows and for aerospace vehicle and propulsion system design.


Recently, multidisciplinary methods have appeared. These methods are used to perform computer simulations of coupled phenomena such as those needed to study wing flutter. Flutter is the motion due to the complex interaction of the instantaneous changes in wing shape due to aerodynamic forces and the instantaneous change in aerodynamic forces due to wing shape. The effect of flutter can be catastrophic if dynamic forces exceed structural limits. However, flutter can, in some cases, be suppressed by a coupled deflection of control surfaces. Simulation of flutter control requires combining fluid dynamics, structural mechanics, and control laws into a time-dependent formulation. One example of simulation of wing flutter control is presented in Fig. 13. In this work<sup>29</sup>, the results were obtained by solving the model structural equations of motion, a simple control law equation, and a nonlinear inviscid form of the fluid dynamic equations (full



AEROELASTICALLY OSCILLATING WING



CONTROL SURFACE ACTIVATED



ACTIVE CONTROL SUPPRESSES OSCILLATION




Fig. 13 Computed pressures on a low aspect ratio wing with and without a control surface activated to suppress flutter.  $M_\infty = 0.9$ ,  $\alpha = 0$ , simulated altitude = 30,000 ft.

potential) in a fully coupled manner. Wing flutter is established with the control surface fixed (upper figure). Next, the control surface is activated (middle) so that its instantaneous deflections reduce and maintain the wing steady (bottom). The various shades of gray on the wing surface represent the levels of instantaneous pressure on the wing.

The example above was a highly approximate treatment for a simple shape. Today's aerospace challenge is to provide integrated, multidisciplinary simulations of aerospace vehicles throughout their mission profiles. This requires the simulation of propulsion systems, complete external aerodynamics, structural response, controls, and chemistry in an integrated manner. It is estimated that these calculations will take tens of hours on computer systems capable of sustaining a trillion floating-point operations per second (teraFLOPS).

In space and Earth sciences, NASA has made significant strides in weather/climate simulation, simulation of stellar/galactic evolution, solid Earth modeling, and analysis of multi-spectral/multi-sensor data. Today's space and Earth sciences challenges are the multidisciplinary modeling and monitoring of the Earth and its global changes, assessments of their potential impact on the future environment of man, better understanding of space phenomena, formation of the universe, and the origins of life. To begin addressing the complex interdisciplinary modeling for such nonlinear problems will require teraFLOPS of computational power and terabytes of memory capability.

In many cases the computational power to solve these problems can be located in ground-based facilities, but in others, such as multi-sensor airborne or satellite systems, the computational power must be located in a flight environment. Highly autonomous space exploration systems, e.g., Mars Rover sample return, will require high-performance and ultra reliable spaceborne computers. These specific grand challenges and others facing the Agency provide a computational imperative for a NASA initiative in High Performance Computing.

The approach to be taken by the HPC Program is to fully exploit highly parallel computing technology by focusing on system architecture, highly parallel algorithms and basic computer science research. One of the program's specific objectives is to develop technologies that will enable a sustained computing rate of a teraFLOPS; approximately 1000 times the sustained performance of an 8-processor Cray YMP. Highly parallel processor technology was chosen because it presents the best alternative to the conventional supercomputer approach. The conventional approach depends on a small number of very fast processors sharing a very fast common memory. This approach will be severely limited in the future by electronic technology limits (speed of light, circuit density, and electronic packaging), processor to memory bandwidth, and high development costs. Highly parallel processors greatly increase parallelism, relax demands on technology, and share development costs with the much larger microcomputer market. Parallel processing research has been ongoing within the NAS program and within other NASA programs for

the past few years. The primary thrust has been to evaluate parallel processing architectures and to develop efficient parallel applications. Results have been encouraging. For example, results have been obtained that show a 16,000-processor CM-2 will out-perform a Cray-2 single processor for certain large computational fluid dynamics applications<sup>30</sup>.

Developing application algorithms and architecture capable of fully utilizing highly parallel concepts is another objective of the program. High-performance, parallel processor testbed facilities will be implemented for developing both applications and systems software. They will be scalable so that initial feasibility can be demonstrated at smaller scales and at reduced cost. The applicability of new parallel processor architectures and algorithms will be demonstrated on selected multidisciplinary research applications. Three areas have been chosen:

- (a) Computational Aerosciences: integrated, multidisciplinary simulation of aerospace vehicles throughout their mission profiles
- (b) Earth and Space Sciences: multidisciplinary modeling and monitoring of the Earth and its global changes and assessment of their impact on the future environment
- (c) Remote Exploration and Experimentation: extended-duration human exploration missions and remote exploration and experimentation

Finally, the basic computer science research infrastructure within NASA will be enhanced both at NASA research centers and at several NASA computationally oriented institutes. Strengthening basic research is crucial to achieving the innovative approaches needed to meet the initiative's goals and exploit newly obtained capabilities.

In summary, highly parallel computers offer potential to leap-frog the current state-of-the-art and accelerate the growth in computer speed. NASA has unique needs in aeronautics and in Earth and space sciences for greatly increased speed. HPCI is aimed at meeting these needs by providing the technology that before the end of the decade can achieve a teraFLOPS or a 1000-fold increase over today's capability. This achievement will enable performance optimization, time and cost reduction, and increased safety and reliability in aerospace vehicles, and establish a powerful resource to maximize scientific benefits of space observations.

## 10. References

1. V. L. Peterson and W. F. Ballhaus, Jr., *History of the Numerical Aerodynamic Simulation Program*, NASA CP-2454 (1987).
2. B. T. Blaylock and F. R. Bailey, "Status and Future Developments of the NAS Processing System Network," *Proceedings of the Third International Conference on Supercomputing*, Vol. II (Boston, Mass., May 15-20, 1988) pp. 9-10.
3. F. R. Bailey and P. Kutler, "NAS-The First Year," *Proceedings of the 16th Congress of the International Council of the Aeronautical Sciences* (Jerusalem, Israel, August 28-September 2, 1988) pp. 1210-1223.
4. D. Choi and C. Levit, "An Implementation of a Distributed Interactive Graphics System for a Supercomputer Environment," *International Journal of Supercomputing Applications*, 1 (Winter 1987) pp. 82-95.
5. S. E. Rogers, P. C. Buning, and F. J. Merritt, "Distributed Interactive Graphics Applications in Computational Fluid Dynamics," *International Journal of Supercomputing Applications*, 1 (Winter 1987) pp. 96-105.
6. T. Lasinski, P. Buning, D. Choi, S. Rogers, G. Bancroft, and F. Merritt, *Flow Visualization of CFD Using Graphics Workstations*, AIAA Paper 87-1180-CP (1987).
7. V. Watson, P. Buning, D. Choi, G. Bancroft, F. Merritt, and S. Rogers, *Use of Computer Graphics for Visualization of Flow Fields*, AIAA Aerospace Engineering Conference and Show (Los Angeles, CA, February 17-19, 1987).
8. R. Magus and H. Yoshihara, "Inviscid Transonic Flow Over Airfoils," *AIAA J.*, 8 (1970) pp. 2157-2162.
9. V. L. Peterson, "Impact of Computers on Aerodynamic Research and Development," *IEEE Proc.* 72 (1984) pp. 68-79.
10. P. Kutler, J. L. Steger, and F. R. Bailey, *Status of Computational Fluid Dynamics in the United States*, AIAA Paper 87-1135-CP (1987).
11. P. Kutler and A. R. Gross, *Progress and Future Directions in Computational Fluid Dynamics*, NASA TM-100091 (1988).

12. D. M. Bushnell, "NASA Research on Viscous Drag Reduction II," in *Proc. Laminar Turbulent Boundary Layers, Energy Sources Technology Conf.*, eds. E. M. Uram and H. E. Weber (New Orleans, LA, February 12-16, 1984).
13. S. K. Robinson, S. J. Kline, and P. R. Spalart, *Spatial Character and Time-Evolution of Coherent Structures in a Numerically Simulated Boundary Layer*, First National Fluid Dynamics Congress (Cincinnati, OH, July 1988).
14. P. R. Spalart, *Direct Simulation of a Turbulent Boundary Layer up to  $Re_{\theta}=1410$* , NASA TM-89407 (1986).
15. V. L. Peterson, J. Kim, T. L. Holst, G. S. Deiwert, D. M. Cooper, A. B. Watson, and F. R. Bailey, "Supercomputer Requirements for Selected Disciplines Important to Aerospace," *Proceedings of the IEEE*, **77** (1989) pp. 1038-1055.
16. T. Holst, *Navier-Stokes Computation Useful in Aircraft Design*, AIAA Paper 90-1800 (1990).
17. P. G. Buning, I. T. Chiu, F. W. Martin, Jr., R. L. Meakin, S. Obayashi, Y. M. Rizk, J. S. Steger, and M. Yarrow, "Flow Field Simulation of the Space Shuttle Vehicle in Ascent," *Proc. Fourth International Conference on Supercomputing and Third World Supercomputer Exhibition*, Vol. II (International Supercomputing Institute, Inc., 1989) pp. 20-28.
18. G. Palmer, *Enhanced Thermochemical Nonequilibrium Computations of Flow Around the Aeroassist Flight Experiment Vehicle*, AIAA Paper 90-1702 (1990).
19. J. N. Moss, G. A. Bird, and K. D. Virendra, *Nonequilibrium Thermal Radiation for an Aeroassist Flight Experiment*, AIAA Paper 88-0081 (1988).
20. C. W. Bauschlicher, S. R. Langhoff, and P. R. Taylor, "Accurate Quantum Mechanical Calculations," *Advances in Chemical Physics*, **77** (1990) p. 103.
21. D. W. Schwenke, "Calculations of Rate Constants for the Three-Body Recombination of  $H_2$  in the Presence of  $H_2$ ," *J. Chem. Phys.* **89** (1988) p. 2076.
22. C. P. Blankenship and R. J. Hayduk, "Potential Supercomputing Needs for Structural Analysis," *Proc. Second International Conference on Supercomputing*, Vol. II (International Supercomputing Institute, Inc., 1987) pp. 180-202.
23. Report of the Presidential Commission on the Space Shuttle Challenger Accident (Washington, DC, June 1986).

24. W. H. Greene, N. F. Knight, Jr., and A. E. Stockwell, *Structural Behavior of the Space Shuttle SRM Tang-Clevis Joint*, NASA TM-89018 (1986).
25. B. F. Smith and R. H. Miller, "Galactic Formation and Evolution," *Proc. Fourth International Conference on Supercomputing and Third World Supercomputer Exhibition*, Vol. II (International Supercomputing Institute, Inc., 1989) pp. 1-6.
26. B. F. Smith and R. H. Miller, "A Computational Approach to Galaxy Formation and Evolution," *Proc. Second International Conference on Supercomputing*, Vol. II (International Supercomputing Institute, Inc., 1987) pp. 207-213.
27. J. B. Pollack, R. M. Haberle, J. Schaeffer, and H. Lee, "Simulation of the General Circulation of the Martian Atmosphere, I. Polar Processes," *J. Geophys. Res.*, **95** (1990) pp. 1447-1473.
28. R. M. Haberle and B. M. Jakosky, "Sublimation and Transport of Water from the North Residual Polar Cap on Mars," *J. Geophys. Res.*, **95** (1990) pp. 1423-1437.
29. G. P. Guruswamy, E. L. Tu, and P. M. Goorjian, *Transonic Aeroelasticity of Wings with Active Control Surfaces*, AIAA Paper 87-0709-CP (1987).
30. C. Levit and D. Jespersen, "Explicit and Implicit Solution of the Navier-Stokes Equations on a Massively Parallel Computer," *Computational Structural Mechanics and Fluid Dynamics—Advances and Trends, Special Issue of Computers and Structures*, **30** (Pergamon Press, 1988).

# Report Documentation Page

1. Report No. NASA TM-102890		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle NASA's Supercomputing Experience				5. Report Date December 1990	
				6. Performing Organization Code	
7. Author(s) F. Ron Bailey				8. Performing Organization Report No. A-91024	
				10. Work Unit No. 505-59	
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035-1000				11. Contract or Grant No.	
				13. Type of Report and Period Covered Technical Memorandum	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546-0001				14. Sponsoring Agency Code	
15. Supplementary Notes Point of Contact: F. Ron Bailey, Ames Research Center, MS 200-4, Moffett Field, CA 94035-1000 (415) 604-5065 or FTS 464-5065  Presented at Singapore Supercomputing Conference 1990, Singapore, December 11-12, 1990.					
16. Abstract  A brief overview of NASA's recent experience in supercomputing is presented from two perspectives: early systems development and advanced supercomputing applications. NASA's role in supercomputing systems development is illustrated by discussion of activities carried out by the Numerical Aerodynamic Simulation Program. Current capabilities in advanced technology applications are illustrated with examples in turbulence physics, aerodynamics, aerothermodynamics, chemistry, and structural mechanics. Capabilities in science applications are illustrated by examples in astrophysics and atmospheric modeling. The paper concludes with a brief comment on the future directions and NASA's new High Performance Computing Program.					
17. Key Words (Suggested by Author(s)) Supercomputing Computational science Numerical simulation			18. Distribution Statement Unclassified-Unlimited  Subject Category - 59		
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 30	22. Price A03		